

Paper EPO-210

A Failure To EXIST: Why Testing for Data Set Existence with the EXIST Function Alone Is Inadequate for Serious Software Development in Asynchronous, Multiuser, and Parallel Processing Environments

Troy Martin Hughes

ABSTRACT

The Base SAS® EXIST function demonstrates the existence (or lack thereof) of a SAS data set. Conditional logic routines commonly rely on EXIST to validate data set existence or absence before subsequent processes can be dynamically executed, circumvented, or terminated based on business logic. In synchronous software design and where data sets cannot be accessed by other processes or users, EXIST is both a sufficient and reliable solution. However, because EXIST captures only a split-second snapshot of the file state, it provides no guarantee of file state persistence. Thus, in asynchronous, multiuser, and parallel processing environments, data set existence can be assessed by one process but instantaneously modified thereafter (by creating or deleting the data set) by a concurrent process, leading to a race condition that causes failure. Due to this vulnerability, most classic implementations of the EXIST function within SAS literature are insufficient for testing data set existence in these more complex environments. This text demonstrates reliable and secure methods to test SAS data set existence prior to performing conditional tasks in asynchronous, multiuser, and parallel processing environments.

INTRODUCTION

To be clear, the EXIST function duly performs as advertised—its credibility and virtue remain unblemished and intact. However, the *implementation* of EXIST can be lacking when SAS practitioners make incorrect assumptions that are translated into faulty business logic. For example, business logic in SAS software often assumes that if a data set exists it is also available—that is, the necessary file lock (either exclusive or shared) will be available. But an existent data set may be unavailable when other users or processes are accessing it concurrently. Thus, methods demonstrated in this text always validate both existence and availability to ensure subsequent processes will not fail due to either missing or locked files.

A second assumption often made in business logic is that EXIST captures a static file state. Yes, EXIST takes a snapshot of file state, but this state can change in a millisecond, causing the results of EXIST to become outdated and incorrect. Where data sets reside in shared libraries accessible by other users or processes, production code that demands robustness and reliability must account for the dynamic nature of data sets and ensure that the file state remains frozen until subsequent tasks have been performed. For example, if software uses the EXIST function to test whether a control table exists, and subsequently prints a report of the data set only when it exists, this logic can fail in a multiuser environment. Immediately after the EXIST statement but before the subsequent REPORT procedure, another user or process could sneak into that infinitesimal interstice and exclusively lock or delete the data set, causing REPORT to fail. Thus, without safeguards to prevent this race condition, conditional business logic that is sufficient for simpler design will fail when faced with multiuser and parallel environments.

Another idiosyncrasy of the implementation of EXIST arises when a negative return value is expected—that is, some action is to be performed when a data set is determined *not* to exist. For example, the existence of a data set is tested, after which the data set is dynamically created if it was found to be missing. However, because a data set that does not exist cannot be locked (to guarantee availability in addition to existence), this can create a failure pattern in which two or more processes simultaneously assess that a data set does not exist and thereafter simultaneously attempt to create it, causing at least one process to fail. This text demonstrates how to overcome this obstacle, using a mutex to lock a pointer to a data set even before the data set has been created. With an assist from the LOCKSAFE macro, the reliable use of EXIST to support conditional logic routines is demonstrated.

While a straightforward invocation of EXIST is sufficient within single-user, serialized, and synchronous program flow, this text expands the capability of EXIST to more complex environments while demonstrating numerous failure patterns that might otherwise thwart would-be adventurers into the world of SAS parallel processing. As the

performance advantages of parallel processing become more widely demonstrated throughout SAS literature and as SAS practitioners continue to shift from monolithic, serialized programming to more efficient and modular parallel processing paradigms, the implementation of some tried-and-true techniques such as the EXIST function will need to be overhauled to ensure the intended functionality is delivered.

EXIST FUNCTION USAGE

The EXIST function tests the existence of a data set (and, implicitly, its respective SAS library) and is often used as a prerequisite to the execution of subsequent conditional logic. The dynamic program flow typically performs one action if the data set exists and another if it does not, operationalized through the SAS macro language. Once data set existence has been demonstrated, subsequent processes typically require either an exclusive or shared lock on the data set. For a background in shared and exclusive locks, please consult a separate text by the author.ⁱ

For example, if the PERM.Control data set exists (i.e., EXIST returns a 1), then its contents might be subsequently printed to the SAS output window, requiring a shared lock on the data set. If the data set does not exist, however, the subsequent PRINT procedure is skipped, thus preventing the runtime error resulting from a missing data set. The following code functions in a single-user environment but can fail when concurrent SAS users or sessions are interacting with the data set:

```
* positive EXIST return value;
%macro print_dsn();
%if %sysfunc(exist(perm.control)) %then %do;
    proc print data=perm.control;
    run;
    %end;
%mend;

%print_dsn;
```

**UNRELIABLE USE
OF EXIST IN
MULTIUSER
ENVIRONMENT**

Conversely, if the PERM.Control data set does not exist (i.e., EXIST returns a 0), the data set might be created through conditional logic, thus ensuring correct functioning of code that follows, and requiring an exclusive lock on PERM.Control while it is being created. Again, the following code uses EXIST as a sufficient control gate in single-user environments but can fail in multiuser environments:

```
* negative EXIST return value;
%macro create_control();
%if %sysfunc(exist(perm.control))=0 %then %do;
    data perm.control;
        length dtg 8 name $10;
    run;
    %end;
%mend;

%macro engine();
%create_control;
* do something with perm.control;
%mend;

%engine;
```

**UNRELIABLE USE
OF EXIST IN
MULTIUSER
ENVIRONMENT**

For many purposes and especially in single-user environments, the previous business logic is sufficient and reliable. However, because EXIST captures only the static file state of a SAS data set, that state could easily be changed by other users or processes. In multiuser environments in which additional users, processes, or SAS sessions may also be interacting with the data set being tested, the EXIST function as previously implemented is insufficient and will lead to heinous and sometimes baffling runtime errors, as demonstrated throughout the following sections.

Three environments precipitate failures in the EXIST function:

- **Multituser** — In a multituser environment, multiple users share a SAS infrastructure, with persistent libraries and data sets accessible to multiple users across some network. Whenever multiple users have access to the same data set, the risk exists that more than one user could attempt to access or perform some action on the same data set concurrently. For this reason, merely operating in a multituser environment adds risks not inherently associated with single-user environments, such as file access collisions and failures of EXIST.
- **Parallel Processing** — In some cases, even a single-user environment can act as (and incur risks similar to) a multituser environment. For example, a SAS practitioner might open two or three SAS sessions on a single laptop at the same time. Software running concurrently on separate SAS sessions acts no differently than multiple users running separate SAS jobs, so again the risk exists that two or more sessions will attempt to access the same data set simultaneously, causing invalid results or failure.
- **Asynchronous Processing** — Asynchronous processing represents a specific subset of parallel processing in which a SAS program executes statements that run at the same time rather than in series as is common in synchronous SAS program flow. For example, by using the SYSTASK statement, one SAS program can spawn multiple batch jobs, each of which executes concurrently in the background. Care must be exercised, however, to ensure that the jobs do not conflict with each other to cause corruption or failure, including unexpected results when EXIST logic is implemented by either the parent or its children.

In all three environments, failure is caused by underlying race conditions in which two or more processes compete for access to the same data set. Race conditions can be difficult to remedy because of the speed at which they occur and, in many cases, the infrequency. Thus, EXIST might work the vast majority of the time even in a multituser environment, but when it does fail, it can have disastrous consequences that lead to the corruption or deletion of a data set. To illustrate the speed of processes, the following code executes only the EXIST function inside a loop so that its average duration can be calculated:

```
%macro speedy;
%let start=%sysfunc(datetime());
%do i=1 %to 10000;
  %put &i;
  %if %sysfunc(exist(perm.control))=0 %then %do;
    %end;
  %end;
%let stop=%sysfunc(datetime());
%put TIME: %sysevalf(&stop-&start) secs;
%mend;

%speedy;
```

Executing over 10,000 times in just 3.5 seconds, each iteration completes in only 0.00035 seconds, or 0.35 milliseconds! Notwithstanding this speed, because other SAS processes are moving just as quickly, other processes can sneak in and alter or gain control of the data set being tested with EXIST. To diminish this possibility, one best practice is to utilize the referenced data set immediately after its existence is tested. For example, in the previous PRINT_DSN macro, the PRINT procedure immediately follows the EXIST function. Similarly, in the CREATE_CONTROL macro, the DATA step immediately follows the EXIST function. However, even this best practice is insufficient to eliminate runtime errors in multituser environments because the interstice can still be invaded by external users or processes. Thus, a more robust and reliable solution is required.

STRESS TESTING EXIST DEMONSTRATES EXIST VULNERABILITY

In some cases, the failure of EXIST in multituser environments does not produce runtime errors but rather invalid results. For example, when a data set is created by setting itself (with the SET statement, presumably to modify the data set while maintaining the original data set name), for a brief shudder, EXIST falsely assesses that the data set does not exist. To demonstrate this anomaly, the following DATA step and CREATE macro should be executed:

```

libname perm 'c:\perm';

data perm.create;
    length char $10;
    char='sas';
run;

%macro create(dsn=, cnt=);
%do i=1 %to &cnt;
    data &dsn;
        set &dsn;
    run;
%end;
%put SYSCC: &syscc;
%mend;

%create(dsn=perm.create, cnt=1000);

```

In a second SAS session, the following TEST_EXIST macro should be executed, which repeatedly tests the existence of the PERM.Create data set:

```

libname perm 'c:\perm';

%macro test_exist(dsn=, cnt=);
%do i=1 %to &cnt;
    %if %sysfunc(exist(&dsn))
        %then %put YES;
    %else %put NO;
%end;
%mend;

%test_exist(dsn=PERM.create, cnt=1000);

```

When the log from TEST_EXIST session is examined, it demonstrates a pattern of predominantly YESes, interspersed with some NOs:

```

YES
YES
YES
NO
YES
YES
YES

```

As it turns out, when SAS builds a data set, it first creates a temporary data set with an .LCK extension. For example, when PERM.Create is being created, the new data set is temporarily named C:\perm\create.sas7bdat.lck, after which this file is renamed C:\perm\create.sas7bdat. Thus, if the EXIST function assesses a data set during this rename process, it will generate inaccurate results that depict that an existent file is missing. This vulnerability occurs in single-user and serialized processing as well, but the threat of external sessions utilizing EXIST is absent so the vulnerability is not exploited.

To demonstrate the frequency of this risk, the TEST_EXIST macro should be modified to the following code, which now counts YESes and NOs (i.e., positive and negative EXIST return codes, respectively) rather than printing them to the log:

```

%macro test_exist(dsn=, cnt=);

```

```

%let yes=0;
%let no=0;
%do i=1 %to &cnt;
    %if %sysfunc(exist(&dsn))
        %then %let yes=%eval(&yes+1);
    %else %let no=%eval(&no+1);
%end;
%put YES: &yes;
%put NO: &no;
%mend;

%test_exist(dsn=PERM.create, cnt=1000);

```

By first executing the CREATE macro in the first session, and by subsequently executing the TEST_EXIST macro in the second session, the results indicate that EXIST failed to detect the PERM.Create data set in 10 out of 1,000 iterations, or 1 percent of the time:

```

%test_exist(dsn=perm.create, cnt=1000);
YES: 990
NO: 10

```

The likelihood of this risk will vary based on processor speed as well as other factors. For example, as a data set increases in size, the ratio of the time spent building the data set (during which the original file name is detectable by EXIST) to the time spent renaming the data set will increase. Thus, the risk of occurrence will be greatest for empty data sets (such as PERM.Create), while very large data sets will very rarely encounter this threat. Notwithstanding frequency of occurrence, the mere fact that EXIST can return invalid results can lead to much more pernicious results, as demonstrated in the next two sections. For example, if rather than merely printing YES or NO or incrementing a counter, other input/output (I/O) functions (such as reading from or writing a data set) are executed based on the results of erroneous EXIST return codes, several failure patterns emerge, both for positive and negative EXIST return codes.

STRESS TESTING POSITIVE EXIST RETURN CODES

This runtime error occurs when EXIST determines that a data exists but, before the next line of code can use the data set, the data set is subsequently deleted or locked by a separate user or process. For example, the following code first tests the existence of a data set and, after validating existence, reads the data set in a DATA step:

```

libname perm 'c:\perm'; * CHANGE BASED ON ENVIRONMENT;

%macro test_positive(dsn=, cnt=);
%do i=1 %to &cnt;
    %put ITERATION: &i;
    %if %sysfunc(exist(&dsn)) %then %do;
        data x;
            set &dsn;
        run;
    %end;
%end;
%mend;

%test_positive(dsn=perm.test, cnt=1000);

```

**UNRELIABLE USE
OF EXIST IN
MULTIUSER
ENVIRONMENT**

The previous code is sufficient in a single-user environment because no other process should be interacting with the PERM.Test data set. However, if other users or processes are simultaneously interacting with the data set, these interactions could change the file state, either by creating the data set or deleting it. To simulate other processes, the following code repeatedly creates and then deletes the PERM.Test data set. To stress test the EXIST function within

the previous business logic, the TEST_POSITIVE macro should be run in one SAS session while the following code (including the DATA step and TEST macro) is run in a separate SAS session:

```
libname perm 'c:\perm'; * CHANGE BASED ON ENVIRONMENT;

data perm.test;
  length char1 $10 num1 8;
run;

%macro test(dsn=, cnt=);
%do i= 1 %to &cnt;
  %put ITERATION: &i;
  data &dsn;
    length iteration dtg 8;
    format iteration 8. dtg datetime17.;
    iteration=&i;
    dtg=%sysfunc(datetime());
  run;
  proc delete data=&dsn;
  run;
%end;
%mend;

%test(dsn=perm.test, cnt=1000);
```

Several different runtime errors should have been produced, although readers may need to increase the number of iterations to demonstrate all error types. The first error occurs when session one (executing the TEST_POSITIVE macro) attempts to access the data set PERM.Test when session two is either creating or deleting the data set, both actions of which require an exclusive file lock. The following output demonstrates this error:

```
ITERATION: 404
ERROR: A lock is not available for PERM.TEST.DATA.

NOTE: The SAS System stopped processing this step because of errors.
WARNING: The data set WORK.X may be incomplete. When this step was stopped there
were 0 observations and 0 variables.
WARNING: Data set WORK.X was not replaced because this step was stopped.
NOTE: DATA statement used (Total process time):
      real time           0.01 seconds
      cpu time            0.01 seconds
```

This illustrates the critical importance of checking for availability (i.e., file lock status) as well as existence whenever a data set must be accessed. The second error type occurs because in the millisecond between the EXIST function and the subsequent DATA step, the DELETE procedure in the second session deletes the data set, thus session one can no longer access it:

```
ITERATION: 433
ERROR: File PERM.TEST.DATA does not exist.

NOTE: The SAS System stopped processing this step because of errors.
WARNING: The data set WORK.X may be incomplete. When this step was stopped there
were 0 observations and 0 variables.
WARNING: Data set WORK.X was not replaced because this step was stopped.
NOTE: DATA statement used (Total process time):
      real time           0.01 seconds
      cpu time            0.01 seconds
```

The third error type occurs when the EXIST function determines that the data set does exist but does so during the DELETE procedure but before the data set has been deleted:

```

ITERATION: 818
ERROR: User does not have appropriate authorization level for file PERM.TEST.DATA.

NOTE: The SAS System stopped processing this step because of errors.
WARNING: The data set WORK.X may be incomplete.  When this step was stopped there
were 0 observations and 0 variables.
WARNING: Data set WORK.X was not replaced because this step was stopped.
NOTE: DATA statement used (Total process time):
      real time           0.01 seconds
      cpu time            0.01 seconds

```

STRESS TESTING NEGATIVE EXIST RETURN CODES

The previous examples performed some action because a data set existed, indicated through the EXIST return code 1. In other cases, an action (e.g., creating a data set) is conditionally performed because a data set does not exist. The following code creates the data set PERM.Test if it does not exist and is sufficient and reliable when executed in a single-user environment:

```

%macro test_negative(dsn=, cnt=);
%do i=1 %to &cnt;
  %put ITERATION: &i;
  %if %sysfunc(exist(&dsn))=0 %then %do;
    data &dsn;
      length char1 $10;
    run;
  %end;
%end;
%mend;

%test_negative(dsn=perm.test, cnt=2);

```

**UNRELIABLE USE
OF EXIST IN
MULTIUSER
ENVIRONMENT**

Thus, the first time the code is executed, it creates the data set PERM.Test (if it didn't already exist) while the loop performs no action the second and subsequent times it is iterated. When the iteration count is increased to 1,000 and TEST_NEGATIVE is executed while the previous TEST macro concurrently executes, several failure patterns again emerge. Note that TEST must be started before TEST_NEGATIVE since TEST_NEGATIVE will otherwise iterate its full 1,000 times before TEST is manually started.

In this sample run, the first observed runtime error is also the most pernicious. Whenever "ERROR: Rename of temporary member" appears in a SAS log, this is an indication that the data set was corrupted during creation and was likely deleted. The log from the SAS session running TEST_NEGATIVE demonstrates this failure pattern:

```

ITERATION: 329

NOTE: Variable char1 is uninitialized.
NOTE: The data set PERM.TEST has 1 observations and 1 variables.
ERROR: Rename of temporary member for PERM.TEST.DATA failed.
      File can be found in c:\perm.
NOTE: DATA statement used (Total process time):
      real time           1.03 seconds
      cpu time            0.12 seconds

ITERATION: 330

```

A second runtime error occurs when EXIST assesses that the data set does not exist but the second SAS session running TEST begins to delete the data set:

ITERATION: 420

ERROR: User does not have appropriate authorization level for file PERM.TEST.DATA.

NOTE: The SAS System stopped processing this step because of errors.

NOTE: DATA statement used (Total process time):

real time 0.01 seconds

cpu time 0.01 seconds

ITERATION: 421

A third runtime error occurs when TEST_NEGATIVE assesses that the data set does not exist, but either the DATA step or DELETE procedure in the TEST macro executes before the DATA step in TEST_NEGATIVE can lock the data set:

ITERATION: 421

ERROR: A lock is not available for PERM.TEST.DATA.

NOTE: The SAS System stopped processing this step because of errors.

WARNING: The data set PERM.TEST was only partially opened and will not be saved.

NOTE: DATA statement used (Total process time):

real time 0.01 seconds

cpu time 0.01 seconds

ITERATION: 422

THE EXISTENCE-AVAILABILITY MATRIX

The existence-availability matrix describes the intersection of the EXIST function return code (either positive or negative) and file lock status (either exclusive or shared). Exclusive locks are read-write and are required whenever a data set is created, modified, or deleted. As the name suggests, only one user or process can gain an exclusive lock on a specific data set at one time. Shared locks conversely are read-only, thus enabling multiple users or processes to gain concurrent access to a data set. For example, PRINT, FREQ, and MEANS procedures can all be run on the same data set at the same time from separate SAS sessions because none of these procedures needs to modify the data. In fact, one SAS session printing a data set will have no awareness that a separate session is simultaneously running the MEANS procedure on the same data set.

Given that EXIST is primarily used to drive conditional logic in which the existence of some data set is assessed before that data set is later used, the relationship between EXIST return codes and the required file lock status can be summarized in three distinct outcomes, represented in Figure 1.

EXIST Return Code	Exclusive Lock	Shared Lock
Positive (EXIST=1)	- modify an existing data set - delete a data set - sort a data set (to same file)	- PROC PRINT or REPORT - PROC MEANS or FREQ - sort a data set (to new file)
Negative (EXIST=0)	create a new data set	N/A

Figure 1. Existence-Availability Matrix with Example Conditional Logic Objectives

For example, if a data set needs to be modified (by setting itself) or deleted, the data set must both exist and have an available exclusive lock, demonstrating that no other user or process is accessing the data set concurrently. This use case is captured in the Positive-Exclusive quadrant of the matrix. When creating a new data set from scratch, conversely, a prerequisite might be that the data set not already exist, which is captured in the Negative-Exclusive quadrant. Finally, the Positive-Shared quadrant captures use cases in which the data set must exist but which do not require exclusive access. Figure 2 demonstrates examples of EXIST functionality in each of these quadrants, with all code demonstrating valid uses of EXIST in single-user environments.

EXIST Return Code	Exclusive Lock	Shared Lock
Positive (EXIST=1)	<pre> %if %sysfunc(exist(&dsn)) %then %do; data &dsn; set &dsn; run; %end; </pre>	<pre> %if %sysfunc(exist(&dsn)) %then %do; proc print data=&dsn; var char1; run; %end; </pre>
Negative (EXIST=0)	<pre> %if %sysfunc(exist(&dsn))=0 %then %do; data &dsn; length char \$10; char='new data'; run; %end; </pre>	N/A

Figure 2. Existence-Availability Matrix with Example EXIST Implementations

While Figures 1 and 2 require the implied exclusive or shared locks, respectively, no mechanism within the Figure 2 examples validates or enforces data set availability. Thus, while testing for existence alone is sufficient in single-user, synchronous, and serialized design, data set availability must additionally be validated within multiuser, asynchronous, and parallel processing design. In the following section, the LOCKSAFE macro is implemented to ensure both data set existence and availability, sufficiently to provide reliable data set locking, effectively freezing the file state while other functions can be performed on the data set.

LOCKSAFE SOLUTION

In a multiuser or parallel environment, the only way to guarantee that a data set exists (or does not exist) for subsequent processes is to lock down its file state, thus preventing external data set access for some duration following the EXIST function. This can be accomplished with the user of mutex semaphores—the business logic and concurrency controls that ensure that no other process attempts to access the data set exclusively at the same time. In a separate text, the author demonstrates reliable file locking with the use of mutex semaphores—techniques that support not only reliable file locking but also reliable EXIST functionality in multiuser environments.ⁱⁱ

To implement LOCKSAFE to support data set existence checking, download the referenced paper and save the LOCKSAFE macro to the following location:

```
C:\perm\locksafe.sas
```

The LOCKSAFE macro relies on semaphore and mutex text files that are created and which indicate whether a data set is locked. The mutexes and semaphores are empty text files that are stored in the MUTEX library, which has a default location of C:\perm\mutex. To change this location, modify the MUTEX library definition within the Locksafe program:

```
libname &mutlib 'c:\perm\mutex'; * CHANGE TO ACTUAL LOCATION
```

The LOCKSAFE invocation should precede the test for existence and, because LOCKSAFE produces a return code LOCKSAFERC, this return code must additionally be tested before use of the EXIST function. For example, the following code demonstrates a modification of the Positive-Exclusive quadrant code from Figure 2, now rendered safe and reliable in multiuser environments:

```
%macro reliable_exist(dsn=);
%locksafe(dsn=&dsn, sec=5, max=600, type=W);
  %if %length(&locksafeRC)=0 %then %do;
    %if %sysfunc(exist(&dsn)) %then %do;
      data &dsn;
      set &dsn;
      run;
    %end;
  %else %put Data Set Does Not Exist!;
  %locksafe(dsn=&dsn, terminate=YES);
  %end;
%else %put Data Set Locked!;
%mend;

%reliable_exist(dsn=perm.control);
```

The EXIST function operates reliably between the two LOCKSAFE invocations because LOCKSAFE effectively freezes file status by preventing other sessions (that are also using LOCKSAFE) from accessing the data set. In other words, between the two LOCKSAFE invocations, EXIST operates as it would in a single-user environment. The second LOCKSAFE invocation is required to terminate the file lock, thus releasing the lock on the PERM.Control data set so it can again be accessed by external SAS sessions if necessary. Without the TERMINATE=YES invocation, the current SAS session would unnecessarily maintain the exclusive lock on the data set and other SAS sessions (that were also using LOCKSAFE) would be indefinitely unable to access the data set.

A benefit of LOCKSAFE is that it can be used to lock data sets that don't yet exist, functionality not available in the now-deprecated SAS LOCK statement, which would return a runtime error when an attempt to lock a nonexistent data set was made. For example, if a process needed to create the PERM.Control data set but the data set did not exist, there exists the possibility that two processes running in separate SAS sessions could each simultaneously assess that the data set did not exist, and each simultaneously attempt to create the data set, leading to failure of one or both processes. Because LOCKSAFE creates a unique pointer to each data set, that pointer can be created before the actual data set exists, thus enabling the lock to be achieved and the data set subsequently created with confidence that no other process is trying to create, modify, or access the same data set.

Implementation and caveats of LOCKSAFE are described in detail in the referenced text. For example, LOCKSAFE must be implemented on all processes that reference a data set or else rogue processes will be able to sneak in and bypass its semaphores and mutexes. Another hiccup can occur when a process locks a data set with the first LOCKSAFE invocation but terminates abruptly before the subsequent TERMINATE=YES invocation. When this occurs, the semaphores and mutexes may need to be manually deleted from the MUTEX library, as described in detail in the LOCKSAFE text.ⁱⁱⁱ

One final caveat of the LOCKSAFE solution supporting EXIST functionality is the lack of shared locking in LOCKSAFE. In its beta release, LOCKSAFE generates only exclusive—not shared—file locks, so the Positive-Shared quadrant of the existence-availability matrix is missing. While an exclusive lock can be generated in lieu of a shared lock, this unnecessarily prevents other users or processes from simultaneously accessing the data set. For example, the following code freezes a file state so that PERM.Control can be printed reliably, even in a multiuser environment:

```
%macro reliable_exist(dsn=);
%locksafe(dsn=&dsn, sec=5, max=600, type=W);
  %if %length(&locksafeRC)=0 %then %do;
    %if %sysfunc(exist(&dsn)) %then %do;
```

```

proc print data=&dsn;
run;
%end;
%else %put Data Set Does Not Exist!;
%locksafe(dsn=&dsn, terminate=YES);
%end;
%else %put Data Set Locked!;
%mend;

%reliable_exist(dsn=perm.control);

```

However, because the PRINT procedure requires only a shared lock (but LOCKSAFE generates an exclusive lock), this implementation unnecessarily prevents separate SAS sessions from concurrently running MEANS, FREQ, or other procedures or processes on PERM.Control that also require only a shared file lock. Future versions of the LOCKSAFE macro could benefit EXIST functionality in the Positive-Shared quadrant by allowing users to specify whether an exclusive or shared lock was required, thus expanding the versatility of LOCKSAFE.

CONCLUSION

The EXIST function has dutifully served the SAS community and while EXIST functions reliably within single-user environments, its implementation within asynchronous, multiuser, and parallel processing environments can lead to puzzling if not disastrous results. Race conditions that can occur between competing processes in separate SAS sessions can cause EXIST to falsely depict that a data set does not exist when it in fact does. Other race conditions can occur when EXIST accurately depicts that a data set does or does not exist but when that data set is immediately created or deleted by a subsequent process in a separate SAS session. This text demonstrates use of the LOCKSAFE macro to overcome these vulnerabilities and to enable EXIST to perform reliably and securely even in more complex, parallel environments.

CONTACT INFORMATION

Your comments and questions are valued and encouraged. Contact the author at:

Name: Troy Martin Hughes

E-mail: troymartinhughes@gmail.com

SAS and all other SAS Institute Inc. product or service names are registered trademarks or trademarks of SAS Institute Inc. in the USA and other countries. ® indicates USA registration.

Other brand and product names are trademarks of their respective companies.

REFERENCES

ⁱ Hughes, Troy Martin. 2014. "From a One-Horse to a One-Stoplight Town: A Base SAS Solution to Preventing Data Access Collisions through the Detection and Deployment of Shared and Exclusive File Locks." Western Users of SAS Software (WUSS) Conference.

ⁱⁱ Hughes, Troy Martin. 2016. "Stress Testing and Supplanting the SAS® LOCK Statement: Implementing Mutex Semaphores To Provide Reliable File Locking in Multiuser Environments To Enable and Synchronize Parallel Processing." Midwest SAS Users Group (MWSUG) Conference.

ⁱⁱⁱ Id.